

I. Introduction

1. “THE truth is out there”:¹ conspiracy theories are all around us. In August 2004, a poll by Zogby International showed that 49 percent of New York City residents, with a margin of error of 3.5 percent, believed that officials of the U.S. government “knew in advance that attacks were planned on or around September 11, 2001, and that they consciously failed to act.”² In a Scripps-Howard Poll in 2006, some 36 percent of respondents assented to the claim that “federal officials either participated in the attacks on the World Trade Center or took no action to stop them.”³ Sixteen percent said that it was either very likely or somewhat likely that “the collapse of the twin towers in New York was aided by explosives secretly planted in the two buildings.”⁴

2. Conspiracy theories can easily be found all over the world. Among sober-minded Canadians, a September 2006 poll found that 22 percent believed that “the attacks on the United States on September 11, 2001 had nothing to do with Osama Bin Laden and were actually a plot by influential Americans.”⁵ In a poll conducted in seven Muslim countries, 78 percent of respondents said that they do not believe the 9/11 attacks were carried out by Arabs.⁶ The most popular account, in these countries, is that 9/11 was the work of the U.S. or Israeli governments.⁷ In China, a bestseller attributes various events (the rise of Hitler, the Asian financial crisis of 1997–1998, and environmental destruction in the developing world) to the Rothschild banking dynasty; the analysis has been read and debated at high levels of business and government, and it appears to have had an effect on discussions about currency policies.⁸ Throughout American history, race-related violence has often been spurred by false rumors, generally pointing to alleged conspiracies by one or another group.⁹

3. What causes such theories to arise and spread? Are they important and perhaps even threatening, or merely trivial and even amusing? What can and should government do about them? We aim here to sketch some psychological and social mechanisms that produce, sustain, and spread these theories; to show that some of them are quite important and should be taken seriously; and to offer suggestions for governmental responses, both as a matter of policy and as a matter of law.

4. Most of the academic literature directly involving conspiracy theories falls into one of two classes: (1) work by analytic philosophers, especially in epistemology and the philosophy of science, that explores a range of issues but mainly asks what counts as a “conspiracy theory” and whether such theories are methodologically suspect;¹⁰ (2) a smattering of work in sociology and Freudian psychology on the causes of conspiracy theorizing.¹¹ We offer some remarks on the conceptual debates here, but we will generally proceed in pragmatic fashion and mostly from the ground up, hewing close to real examples and the policy problems they pose. To illuminate issues of policy, we draw upon literatures in social psychology, economics, and other disciplines concerning informational cascades, the spread of rumors, and the epistemology of groups and social networks. We adapt the insights of these literatures by focusing on the features of false and harmful conspiracy theories that make them distinct from, and sometimes more damaging than, other false and harmful beliefs.

5. Our running example involves conspiracy theories relating to terrorism, especially theories that arose from and post-date the 9/11 attacks. Terrorism-related theories are hardly the only ones of interest, but they provide a crucial testing ground for the significance, causes, and policy implications of widespread conspiracy theorizing. As we shall see, an understanding of conspiracy theories illuminates

the spread of information and beliefs more generally. We shall also see, however, that because of their special characteristics, conspiracy theories pose unique challenges.

6. Section I explores some definitional issues and lays out some of the mechanisms that produce conspiracy theories and theorists. We begin by narrowing our focus to conspiracy theories that are false, harmful, and unjustified (in the epistemological sense), and by discussing different understandings of the nature of such conspiracy theories and different accounts of the kinds of errors made by those who hold them. Our primary claim is that those who hold conspiracy theories of this distinctive sort typically do so not as a result of a mental illness of any kind, or of simple irrationality, but as a result of a “crippled epistemology,” in the form of a sharply limited number of (relevant) informational sources. In that sense, acceptance of such theories may not be irrational or unjustified from the standpoint of those who adhere to them within epistemologically isolated groups or networks,¹² although they are unjustified relative to the information available in the wider society, especially if it is an open one. There is a close connection, we suggest, between our claim on this count and the empirical association between terrorist behavior and an absence of civil rights and civil liberties.¹³ When civil rights and civil liberties are absent, people lack multiple information sources, and they are more likely to have reason to accept conspiracy theories.

7. Section II discusses government responses and legal issues. We address several dilemmas of governmental response to false, harmful, and unjustified conspiracy theories. Conspiracy theories turn out to be unusually hard to undermine or dislodge; they have a self-sealing quality, rendering them particularly immune to challenge. Our principal claim here involves the potential value of *cognitive infiltration of extremist groups*, designed to introduce informational diversity into such groups and to expose indefensible conspiracy theories as such.

II. Definitions and Mechanisms

A. Definitional Notes

8. There has been much discussion of what, exactly, counts as a conspiracy theory, and about what, if anything, is wrong with those who hold one. Of course it would be valuable to specify necessary and sufficient conditions for such theories, in a way that would make it possible to make relevant distinctions. However, the various views that people label “conspiracy theories” may well relate to each other through a family-resemblance structure, such that necessary and sufficient conditions cannot be given even in principle.

9. We bracket the most difficult conceptual questions here and suggest, pragmatically, that a conspiracy theory can generally be counted as such if it is *an effort to explain some event or practice by reference to the machinations of powerful people, who attempt to conceal their role (at least until their aims are accomplished)*. While many conspiracy theories involve people who are not especially powerful (friends, neighborhoods, fellow employees, family members, and so forth), this account is the most useful for our particular purposes, and it seems to capture the essence of the most prominent and influential conspiracy theories about public affairs. Consider, for example, the view that the Central Intelligence Agency was responsible for the assassination of President John F. Kennedy; that doctors deliberately manufactured the AIDS virus; that the 1996 crash of TWA flight 800 was caused by a U.S. military missile; that the theory of global warming is a deliberate fraud; that the Trilateral Commission is responsible for

important movements of the international economy; that Martin Luther King, Jr. was killed by federal agents; that the plane crash that killed Democrat Paul Wellstone was engineered by Republican politicians; that the moon landing was staged and never actually occurred; that the Rothschilds and other Jewish bankers are responsible for the deaths of presidents and for economic distress in Asian nations; and that the Great Depression was a result of a plot by wealthy people to reduce the wages of workers.¹⁴

10. Of course some conspiracy theories have turned out to be true, and under our definition, they do not cease to be conspiracy theories for that reason.¹⁵ The Watergate hotel room used by Democratic National Committee was, in fact, bugged by Republican officials, operating at the behest of the White House. In the 1950s, the Central Intelligence Agency did, in fact, administer LSD and related drugs under Project MKULTRA, in an effort to investigate the possibility of “mind control.” Operation Northwoods, a rumored plan by the Department of Defense to simulate acts of terrorism and to blame them on Cuba, really was proposed by high-level officials (though the plan never went into effect).¹⁶ Our focus throughout is on demonstrably false conspiracy theories, such as the various 9/11 conspiracy theories, not ones that are true or whose truth is undetermined. Our ultimate goal is to explore how public officials might undermine such theories, and as a general rule, true accounts should not be undermined.¹⁷

11. Within the set of false conspiracy theories, we also limit our focus to potentially harmful theories. Consider the false conspiracy theory, held by many of the younger members of our society, that the mysterious “Santa Claus” distributes presents around the world on Christmas Eve. This theory turns out to be false, but is itself instilled through a widespread conspiracy of the powerful—parents—who conceal their role in the whole affair. It is an open question whether most conspiracy theories are equally benign; we will suggest that some are not benign at all.

12. Under this account, conspiracy theories are a subset of the larger category of false beliefs, and also of the somewhat smaller category of beliefs that are both false and harmful. Consider, for example, the beliefs that prolonged exposure to sunlight is actually healthy, that cigarette smoking does not cause cancer, and that climate change is neither occurring nor likely to occur. These beliefs are (in our view) both false and dangerous, but as stated, they need not depend on, or posit, any kind of conspiracy theory. We shall see that the mechanisms that account for conspiracy theories overlap with those that account for false and dangerous beliefs of all sorts, including those that fuel anger and hatred.¹⁸ But as we shall also see, conspiracy theories have distinctive features, above all because of their self-sealing quality; the very arguments that give rise to them, and account for their plausibility, make it more difficult for outsiders to rebut or even to question them.

13. Conspiracy theories often attribute extraordinary powers to certain agents—to plan, to control others, to maintain secrets, and so forth. Those who believe that those agents have such powers are especially unlikely to give respectful attention to debunkers, who may, after all, be agents or dupes of those who are responsible for the conspiracy in the first instance. It is comparatively easier for government to dispel false and dangerous beliefs that rest, not on a self-sealing conspiracy theory, but on simple misinformation or on an apparent or actual social consensus that is fragile and easily “tipped.”¹⁹ The most direct governmental technique for dispelling false (and also harmful) beliefs—providing credible public information—does not work, in any straightforward way, for conspiracy theories. This

extra resistance to correction through simple techniques is what makes conspiracy theories distinctively worrisome.

14. A further question about conspiracy theories—whether true or false, harmful or benign—is whether they are justified. Justification and truth are different issues, which is why pointing out that some conspiracy theories are true does not show that it is rational to believe in those theories. A true belief may be unjustified, and a justified belief may be untrue. I may believe, correctly, that there are fires within the earth’s core, but if I believe that because the god Vulcan revealed it to me in a dream, my belief is unwarranted. Conversely, the false belief in Santa Claus is justified, because children generally have good reason to believe what their parents tell them and follow a sensible heuristic (“if my parents say it, it is probably true”); when children realize that Santa is the product of a widespread conspiracy among parents, they have a justified and true belief that a conspiracy has been at work.

15. Our final narrowing condition is that we are concerned only with (the many) conspiracy theories that are false, harmful, and unjustified (not in the sense of being irrationally held by those individuals who hold them, but from the standpoint of the information available in the society as a whole). When and under what conditions are conspiracy theories unjustified? Here there are competing accounts and many controversies, in epistemology and analytic philosophy.²⁰ We need not opt for only one of these accounts, because they are not mutually exclusive; each accounts for part of the terrain.

16. Karl Popper famously argued that conspiracy theories overlook the pervasive unintended consequences of political and social action; they assume that all consequences must have been intended by someone.²¹ Many social effects, including large movements in the economy, occur as a result of the acts and omissions of many people, none of whom intended to cause those effects. The appeal of some conspiracy theories, then, lies in the attribution of otherwise inexplicable events to intentional action,²² and to an unwillingness to accept the possibility that significant adverse consequences may be a product of invisible hand mechanisms (such as market forces or evolutionary pressures) or of simple chance,²³ rather than of anyone’s plans.²⁴

17. Popper captures an important feature of some conspiracy theories. There is a pervasive human tendency to think that effects are caused by intentional action, especially by those who stand to benefit (the “cui bono?” maxim), and for this reason conspiracy theories have considerable but unwarranted appeal.²⁵ On one reading of Popper’s account, those who accept conspiracy theories are following a sensible heuristic, to the effect that consequences are intended; that heuristic often works well, but it also produces systematic errors, especially in the context of outcomes that are a product of complex interactions among numerous people. More broadly, Popper is picking up on a general fact about human psychology, which is that most people do not like to believe that significant events were caused by bad (or good) luck, and much prefer simpler causal stories.²⁶ In particular, human “minds protest against chaos,” and people seek to extract a meaning from a bewildering event or situation,²⁷ a meaning that a conspiracy may well supply.

18. Note, however, that the domain of Popper’s explanation is quite limited. Many conspiracy theories, including those involving political assassinations and the attacks of 9/11, point to events that are indeed the result of intentional action, and the conspiracy theorists go wrong not by positing

intentional actors, but by misidentifying them. (The theory that Al-Qaeda was responsible for 9/11 is thus a justified and true conspiracy theory.)

19. Conspiracy theories that posit machinations by government officials typically overestimate the competence and discretion of officials and bureaucracies, who are assumed to be able to make and carry out sophisticated secret plans, despite abundant evidence that in open societies government action does not usually remain secret for very long.²⁸ Consider all the work that must be done to hide and to cover up the government's role in producing a terrorist attack on its own territory, or in arranging to kill political opponents.

20. In a closed society, secrets are far easier to keep, and distrust of official accounts makes a great deal of sense. In such societies, conspiracy theories are both more likely to be true and harder to show to be false in light of available information. But when the press is free, and when checks and balances are in force, it is harder for government to keep nefarious conspiracies hidden for long. These points do not mean that it is logically impossible, even in free societies, that conspiracy theories are true; sometimes they are. But it does mean that institutional checks make it less likely, in such societies, that powerful groups can keep dark secrets for extended periods, at least if those secrets involve illegal or nefarious conduct. Of course conspiracy theories are widespread even in open societies, including the United States, the United Kingdom, and France; the only point is that such theories are less likely to be either true or justified in such societies.

21. An especially useful account suggests that what makes unjustified conspiracy theories unjustified is that those who accept them must also accept a kind of spreading distrust of all knowledge-producing institutions, in a way that makes it difficult to believe anything at all.²⁹ To think, for example, that U.S. government officials destroyed the World Trade Center and then covered their tracks requires an ever-widening conspiracy theory, in which the 9/11 Commission, congressional leaders, the FBI, and the media were either participants in or, at best, dupes of the conspiracy. But anyone who believed that would undercut the grounds for many of their other beliefs, which are warranted only by trust in the knowledge-producing institutions created by government and society. As Robert Anton Wilson notes of the conspiracy theories advanced by Holocaust deniers, "a conspiracy that can deceive us about 6,000,000 deaths can deceive us about anything, and [then] it takes a great leap of faith for Holocaust Revisionists to believe World War II happened at all, or that Franklin Roosevelt did serve as President from 1933 to 1945, or that Marilyn Monroe was more 'real' than King Kong or Donald Duck."³⁰

22. This is not, and is not be intended to be, a general claim that conspiracy theories are unjustified or unwarranted in all imaginable situations or societies. Much depends on the background state of knowledge-producing institutions. If those institutions are generally trustworthy, in part because they are embedded in an open society with a well-functioning marketplace of ideas and free flow of information, and if it is difficult to dupe many diverse institutions simultaneously (as the 9/11 conspiracy theories require), then conspiracy theories will usually be unjustified. On the other hand, individuals in societies with systematically malfunctioning or skewed institutions of knowledge—say, individuals who live in an authoritarian regime lacking a free press—may have good reason to distrust all or most of the official denials they hear. For these individuals, conspiracy theories will more often be warranted, whether or not true.

23. Likewise, individuals embedded in isolated groups or small, self-enclosed networks who are exposed only to skewed information will more often hold conspiracy theories that are justified, relative to their limited informational environment.³¹ Holocaust denials might themselves be considered in this light. When epistemologically isolated groups operate within a society that is both wider and more open, their theories may be unjustified from the standpoint of the wider society but justified from the standpoint of the individual or group. In these situations, the problem for the wider society is to breach the informational isolation of the small group or network (a problem we discuss below).

24. On our account, a central feature of conspiracy theories is that they are extremely resistant to correction, certainly through direct denials or counterspeech by government officials; apparently contrary evidence can usually be shown to be a product of the conspiracy itself. Conspiracy theories often display the characteristic features of a “degenerating research program”³² in which contrary evidence is explained away by adding epicycles and resisting falsification of key tenets.³³ Some epistemologists argue that this resistance to falsification is not objectionable if one also believes that there are conspirators deliberately attempting to plant evidence that would falsify the conspiracy theory, thereby covering their tracks.³⁴ However that may be as a philosophical matter, the self-sealing quality of conspiracy theories creates severe practical problems for government; direct attempts to dispel the theory can usually be folded into the theory itself, as just one more ploy by powerful conspiracy members.

25. So far we have discussed some epistemological features of conspiracy theories, in the abstract; narrowed our focus to conspiracy theories that are false, harmful, and unjustified from the standpoint of the wider society (although they may be justified from the standpoint of individuals, given the information they have received, or within the closed epistemological network of the conspiracy theorists); and suggested institutional grounds for thinking that in a free and open society, there is usually good reason to believe that most conspiracy theories will lack adequate justification. We now turn to the sociology of conspiracy theorizing, examining the mechanisms by which such theories arise and expand.

B. HOW CONSPIRACY THEORIES ARISE AND SPREAD

i. Crippled Epistemologies

26. Why do people accept conspiracy theories that turn out to be false and for which the evidence is weak or even nonexistent? It is tempting to answer in terms of individual pathology, either literal or metaphorical.³⁵ Perhaps conspiracy theories are a product of mental illness, such as paranoia or narcissism, or of similar conditions. And surely some people who accept conspiracy theories are mentally ill and subject to delusions.³⁶ But we have seen that in many communities and even nations, such theories are widely held. It is not plausible to suggest that all or most members of those communities are afflicted by mental illness; the metaphor of mental illness itself obscures more than it clarifies. The most important conspiracy theories are hardly limited to those who suffer from any kind of pathology.

27. For our purposes, the most useful way to understand the pervasiveness of conspiracy theories is to examine how people acquire their beliefs.³⁷ For most of what they believe that they know, human beings lack personal or direct information; they must rely on what other people think. In some domains, people suffer from a “crippled epistemology,” in the sense that they know very few things, and what

they know is wrong.³⁸ Many extremists fall in this category; their extremism stems not from irrationality, but from the fact that they have little (relevant) information, and their extremist views are supported by what little they know.³⁹ Conspiracy theorizing often has the same feature. Those who believe that Israel was responsible for the attacks of 9/11, or that the Central Intelligence Agency killed President Kennedy, may well be responding quite rationally to the informational signals that they receive; in this sense, those beliefs may well be justified from the standpoint of the individuals who hold them, even if they are preposterous in light of the information available in the wider society.

28. Consider here the suggestive claim that terrorism is more likely to arise in nations that lack civil rights and civil liberties.⁴⁰ If this is so, it might be because terrorism is not abstract violence but an extreme form of political protest, and when people lack the usual outlets for registering their protest, they might resort to violence.⁴¹ But consider another possibility: when civil rights and civil liberties are restricted, little information is available, and what comes from government cannot be trusted. If the most trustworthy or least untrustworthy information justifies conspiracy theories and (therefore) extremism, and (therefore?) violence, then terrorism is more likely to arise.

ii. Rumors and Speculation

29. Of course it is necessary to specify how, exactly, conspiracy theories begin. Some such theories seem to bubble up spontaneously, appearing roughly simultaneously in many different social networks; others are initiated and spread, quite intentionally, by conspiracy entrepreneurs who profit directly or indirectly from propagating their theories. One example in the latter category is the author of the Chinese bestseller mentioned above;⁴² another is the French author Thierry Meyssan, whose book “9/11: The Big Lie” became a bestseller and a sensation for its claims that the Pentagon explosion on 9/11 was caused by a missile, fired as the opening salvo of a coup d’etat by the military-industrial complex, rather than by American Airlines Flight 77.⁴³

30. Some conspiracy entrepreneurs are entirely sincere. Others are interested in money or power or in using the conspiracy theory to achieve some general social goal. In the context of the AIDS virus, for example, a diverse set of people initiated rumors, many involving conspiracies, and in view of the confusion and fear surrounding that virus, several of those rumors spread widely.⁴⁴ But even for conspiracy theories put about by conspiracy entrepreneurs, the key question is why some theories take hold while many more do not, and vanish into obscurity.

31. Whenever a bad event has occurred, rumors and speculation are inevitable. Most people are not able to know, on the basis of personal or direct knowledge, why an airplane crashed, or why a leader was assassinated, or why a terrorist attack succeeded, or why many people stayed in an area despite what turned out to be an imminent natural disaster. In the aftermath of such an event, numerous speculations will be offered, and some of them will likely point to some kind of conspiracy. To some people, those speculations will seem plausible, perhaps because they provide a suitable outlet for outrage and blame, perhaps because the speculation fits well with other deeply rooted beliefs that they hold. Terrible events produce outrage, and when people are outraged, they are all the more likely to seek causes that justify their emotional states, and also to attribute those events to intentional action.⁴⁵ Conspiracy theories, like rumors, may simultaneously relieve “a primary emotional urge” and offer an explanation, to those who accept the theory, of why they feel as they do; the theory “rationalizes while it relieves.”⁴⁶

32. In addition, antecedent beliefs are a key to the success or failure of conspiracy theories. Within the United States, some people would find it impossibly jarring to think that the Central Intelligence Agency (CIA) was responsible for the assassination of a civil rights leader; that thought would unsettle too many of their other judgments. Others would find those other judgments strongly supported, even confirmed, by the suggestion that the CIA was responsible for such an assassination. Compare the case of terrorist attacks. For most Americans, a claim that the United States government attacked its own citizens, or condoned such attacks, would make it impossible to hold onto a wide range of other judgments. Clearly this point does not hold for many people in Islamic nations, for whom it is far from jarring to believe that responsibility lies with the United States (or Israel).

iii. Conspiracy Cascades: The Role of Information

33. To see how informational cascades work, imagine a group of people who are trying to assign responsibility for some loss of life. Assume that the group members are announcing their views in sequence. Each member attends, reasonably enough, to the judgments of others. Andrews is the first to speak. He suggests that the event was caused by a conspiracy of powerful people. Barnes now knows Andrews's judgment; she should certainly go along with Andrew's account if she agrees independently with him. But if her independent judgment is otherwise, she would—if she trusts Andrews no more and no less than she trusts herself—be indifferent about what to do, and she might simply flip a coin.

34. Now turn to a third person, Charleton. Suppose that both Andrews and Barnes have endorsed the conspiracy theory, but that Charleton's own view, based on limited information, suggests that they are probably wrong. In that event, Charleton might well ignore what he knows and follow Andrews and Barnes. It is likely, after all, that both Andrews and Barnes had evidence for their conclusion, and unless Charleton thinks that his own information is better than theirs, he should follow their lead. If he does, Charleton is in a cascade. Of course Charleton will resist if he has sufficient grounds to think that Andrews and Barnes are being foolish. But if he lacks those grounds, he is likely to go along with them. This may happen even if Andrews initially speculated in a way that does not fit the facts. That initial speculation, in this example, can start a process by which a number of people are led to participate in a cascade, accepting a conspiracy theory whose factual foundations are fragile.

35. Of course the example is highly stylized; conspiracy cascades arise through more complex processes, in which diverse thresholds are crucial. In a standard pattern, the conspiracy theory is initially accepted by people with low thresholds for its acceptance. Perhaps the theory is limited, in its acceptance, to those with such thresholds. But sometimes the informational pressure builds, to the point where many people, with somewhat higher thresholds, begin to accept the theory too. And when many people hold that belief, those with even higher thresholds may come to accept the theory, leading to widespread acceptance of falsehoods. In theory, a conspiracy theory might be justifiably held by many even though it is false and harmful, and even though only a few early movers suggested a strong commitment to it. As a real-world example of a conspiracy cascade, consider the existence of certain judgments about the origins and causes of AIDS, with some groups believing, implausibly, that the virus was produced in government laboratories.⁴⁷ These and other views about AIDS are a product of social interactions, and in particular of cascade effects.

iv. Conspiracy Cascades: The Role of Reputation

36. Conspiracy theories do not take hold only because of information. Sometimes people profess belief in a conspiracy theory, or at least suppress their doubts, because they seek to curry favor.⁴⁸ Reputational pressures help account for conspiracy theories, and they feed conspiracy cascades. In a reputational cascade, people think that they know what is right, or what is likely to be right, but they nonetheless go along with the crowd in order to maintain the good opinion of others.

37. Suppose that Albert suggests that the Central Intelligence Agency was responsible for the assassination of President Kennedy, and that Barbara concurs with Albert, not because she actually thinks that Albert is right, but because she does not wish to seem, to Albert, to be some kind of dupe. It should be easy to see how this process might generate a cascade. Once Albert, Barbara, and Cynthia offer a united front on the issue, their friend David might be reluctant to contradict them even if he believes that they are wrong. In real-world conspiracy theories, reputational pressures often play a large role, leading people to squelch their own doubts in order to avoid social sanctions.

v. Conspiracy Cascades: The Role of Availability

38. Informational and reputational cascades can occur without any particular triggering event. But a distinctive kind of cascade arises when such an event is highly salient or cognitively “available.”⁴⁹ In the context of many risks, such as those associated with terrorism, nuclear power, and abandoned hazardous waste dumps, a particular event initiates a cascade, and it stands as a trigger or a symbol justifying public concern, whether or not that concern is warranted.⁵⁰ Often political actors, both self-interested and altruistic, work hard to produce such cascades.

39. Conspiracy theories are often driven through the same mechanisms. A particular event becomes available,⁵¹ and conspiracy theories are invoked both in explaining it and using it as a symbol for broader social forces and large narratives about political life, casting doubt on accepted wisdom in many domains. Within certain nations and groups, the claim that the United States or Israel was responsible for the attacks of 9/11 fits well within a general narrative about who is the aggressor, and the liar, in a series of disputes—and the view that Al Qaeda was responsible raises questions about that same narrative.

vi. Conspiracy Cascades: The Role of Emotions

40. Thus far our account has been purely cognitive: conspiracy theories circulate in the same way that other beliefs circulate, as people give weight to the views of others and attend to their own reputations. But it is clear that affective factors, and not mere information, play a large role in the circulation of rumors of all kinds. Many rumors persist and spread because they serve to justify or to rationalize an antecedent emotional state produced by some important event, such as a disaster or a war.⁵² When people are especially angry or fearful, they are more likely to focus on particular sorts of rumors and to spread them to others. And when rumors trigger intense feelings, they are far more likely to be circulated.

41. Experimental evidence strongly supports this speculation in the analogous context of “urban legends.”⁵³ When urban legends—involving, for example, a decapitated motorcycle rider, a rat in a soda bottle, or cat food mislabelled as tuna—are devised so as to trigger strong emotions (such as disgust), people are more likely to pass them along. Perhaps the most revealing of these experiments involved actual spreading of urban legends on the Internet.⁵⁴ The conclusion is that in the marketplace of ideas,

“emotional selection” plays a significant role, and it helps to explain such diverse phenomena as moral panics about deviant behavior, hysteria about child abuse, and media attention to relatively small sources of risk such as road rage and “flesh eating bacteria.”⁵⁵ A particular problem involves “emotional snowballing”—runaway selection for emotional content rather than for information.⁵⁶

42. The applications to conspiracy theories should not be obscure. When a terrible event has occurred, acceptance of such theories may justify or rationalize the affective state produced by that event; consider conspiracy theories in response to political assassinations.⁵⁷ In addition, such theories typically involve accounts, or rumors, that create intense emotions, such as indignation, thus producing a kind of emotional selection that will spread beliefs from one person to another.⁵⁸ Of course evidence matters, and so long as there is some kind of process for meeting falsehoods with truth, mistaken beliefs can be corrected. But sometimes the conditions for correction are not present.

vii. Group Polarization

42. There are clear links between cascades and the well-established phenomenon of group polarization, by which members of a deliberating group typically end up in a more extreme position in line with their tendencies before deliberation began.⁵⁹ Group polarization has been found in hundreds of studies involving over a dozen countries.⁶⁰ Belief in conspiracy theories is often fueled by group polarization.⁶¹

43. Consider, as the clearest experimental example, the finding that those who disapprove of the United States, and are suspicious of its intentions, will increase their disapproval and suspicion if they exchange points of view. There is specific evidence of this phenomenon among citizens of France: with respect to foreign aid, they trust the United States a great deal less, and suspect its intentions a great deal more, after they talk with one another.⁶² It should be easy to see how similar effects could occur for conspiracy theories. Those who tend to think that Israel was responsible for the attacks of 9/11, and who speak with one another, will end up with a greater commitment to that belief.⁶³ One result of group polarization is that different groups may end up with radically different attitudes toward conspiracy theories in general and in particular. Speaking with like-minded others, some people may come to find such a theory irresistible and others may come to find it preposterous.⁶⁴

44. Group polarization occurs for reasons that parallel the mechanisms that produce cascades.⁶⁵ Informational influences play a large role. In any group with some initial inclination, the views of most people in the group will inevitably be skewed in the direction of that inclination. As a result of hearing the various arguments, social interactions will lead people toward a more extreme point in line with what group members initially believed. Reputational factors matter as well. Once people hear what others believe, some will adjust their positions at least slightly in the direction of the dominant position. Relatedly, group polarization can occur through positional jockeying; if, for example, several members of a group want to be the second-most-extreme supporter of the view held in common by the group, the distribution of views within the group can shift to become more extreme overall.

45. For purposes of understanding the spread of conspiracy theories, it is especially important to note that group polarization is particularly likely, and particularly pronounced, when people have a shared sense of identity and are connected by bonds of solidarity.⁶⁶ These are circumstances in which

arguments by outsiders, unconnected with the group, will lack much credibility, and fail to have much of an effect in reducing polarization.

viii. Selection Effects

46. A crippled epistemology can arise not only from informational and reputational dynamics within a given group, but also from self-selection of members into and out of groups with extreme views.⁶⁷ Once polarization occurs or cascades arise, and the group's median view begins to move in a certain direction, doubters and halfway-believers will tend to depart while intense believers remain. The overall size of the group may shrink, but the group may also pick up new believers who are even more committed, and in any event the remaining members will, by self-selection, display more fanaticism. Group members may engage in a kind of double-think, segregating themselves, in a physical or informational sense, in order to protect their beliefs from challenge by outsiders.⁶⁸ Even if the rank and file cannot coherently do this, group leaders may enforce segregation in order to insulate the rank and file from information or arguments that would undermine the leaders' hold on the group. As a result, group polarization will likely intensify.

47. Members of informationally and socially isolated groups become increasingly distrustful or suspicious of the motives of others or of the larger society, falling into a "sinister attribution error."⁶⁹ This error occurs when people feel that they are under pervasive scrutiny, and hence they attribute personalistic motives to outsiders and overestimate the amount of attention they receive. Benign actions that happen to disadvantage the group are taken as purposeful plots, intended to harm.⁷⁰ Although these conditions resemble individual-level pathologies such as paranoid cognition, they arise from the social and informational structure of the group, especially those operating in enclosed or closely knit networks, and are not usefully understood as a form of mental illness, not even in a metaphorical sense.

III. Governmental Responses

What can the government do about conspiracy theories, and what should it do?

48. (1) Government might ban "conspiracy theories", somehow defined. (2) Government might impose some kind of tax, financial or otherwise, on those who disseminate such theories. (3) Government might itself engage in counterspeech, marshaling arguments to discredit conspiracy theories. (4) Government might formally hire credible private parties to engage in counterspeech. (5) Government might engage in informal communication with such parties, encouraging them to help. Each instrument has a distinctive set of potential effects, or costs and benefits, and each will have a place under imaginable conditions. Our main policy claim here is that government should engage in cognitive infiltration of the groups that produce conspiracy theories, which involves a mix of (3), (4), and (5).

49. The first-line response to conspiracy theories is to maintain an open society, in which those who might be tempted to subscribe to such theories are unlikely to distrust all knowledge-creating institutions, and are exposed to evidence and corrections. Nongovernmental organizations, including the media, can and do work hard to respond to such theories. As an ambitious example, consider an Internet site, www.snopes.com, which researches rumors and conspiracy theories and reports on their truth or

falsity. (Another is www.counterknowledge.com). For those concerned about the proliferation of conspiracy theorizing on the Internet, this site provides a reliable and helpful reality check. It would be easy to imagine other ventures, small and large, in this vein, for the Internet provides not only a mechanism by which to spread conspiracy theories, but also a range of corrective tools. The more general point is that in free societies, conspiracy theories are generally dislodged by the media and other non-governmental actors.

50. But we have seen that even in open societies, conspiracy theories have some traction.⁷¹ The Internet itself has ambiguous effects; it is no panacea, even if it is entirely free. Reduced information costs may make it easier for private monitors to rebut conspiracy theories,⁷² yet by the same token that very reduction in information costs also makes it easier for conspiracy theorists to generate and spread their theories in the first place. The overall effect of new technology is unclear, as is the ability of nongovernmental monitors to police conspiracy theories. In part because this is so, an official response will sometimes be essential, at least in important cases. And in all cases, the government must necessarily decide whether to deny, confirm, or ignore conspiracy theories that come to the attention of officials.

51. Here we suggest two concrete ideas for government officials attempting to fashion a response to such theories. First, responding to more rather than fewer conspiracy theories has a kind of synergy benefit: it reduces the legitimating effect of responding to any one of them, because it dilutes the contrast with un rebutted theories. Second, we suggest a distinctive tactic for breaking up the hard core of extremists who supply conspiracy theories: cognitive infiltration of extremist groups, whereby government agents or their allies (acting either virtually or in real space, and either openly or anonymously) will undermine the crippled epistemology of believers by planting doubts about the theories and stylized facts that circulate within such groups, thereby introducing beneficial cognitive diversity.

52. Throughout, we assume a well-motivated government that aims to eliminate conspiracy theories, or draw their poison, if and only if social welfare is improved by doing so. (We do not offer a particular account of social welfare, taking the term instead as a placeholder for the right account.) This is a standard assumption in policy analysis, although real-world governments can themselves be purveyors of conspiracy theories, as when the Bush administration suggested that Saddam Hussein had conspired with Al Qaeda to support the 9/11 attacks.⁷³

A. Are Conspiracy Theories Consequential?

53. One line of thinking denies that conspiracy theories matter.⁷⁴ There are several possible reasons to think so. First, conspiracy theories may be held by only a tiny fraction of the relevant population. Second, even if a particular conspiracy theory is widely held in the sense that many people will confess to it when polled, conspiracy theories may be held as “quasi-beliefs”—beliefs that are not costly and possibly even fun to hold, like a belief in UFOs, and that do not form a premise for action.⁷⁵ Perhaps those who seem to accept such theories have “soft” beliefs, in a way that leads them generally to keep quiet, and rarely to act on what they tend to think.

54. It is true that many people do not in fact take any action on the basis of their mistaken beliefs. But in the racial context, a belief in conspiracies has often played a significant role in producing

violence; conspiracy theories have had large effects on behavior.⁷⁶ And even if only a small fraction of adherents to a particular conspiracy theory act on the basis of their beliefs, that small fraction may be enough to cause serious harms. Consider the Oklahoma City bombing, whose perpetrators shared a complex of conspiratorial beliefs about the federal government. Many who shared their beliefs did not act on them, but a few actors did, with terrifying consequences. James Fearon and others argue that technological change has driven down the costs of delivering attacks with weapons of mass destruction, to the point where even a small group can pose a significant threat.⁷⁷ If so, and if only a tiny fraction of believers act on their beliefs, then as the total population with conspiratorial beliefs grows, it becomes nearly inevitable that action will ensue.

55. In other, perhaps more common, cases the conspiracy theory will be of a different nature and will not directly produce such action. However, such theories can still have pernicious effects from the government's point of view, either by inducing unjustifiably widespread public skepticism about the government's assertions, or by dampening public mobilization and participation in government-led efforts, or both.

B. Dilemmas and Responses

56. Imagine a government facing a population in which a particular conspiracy theory is becoming widespread. We will identify two basic dilemmas that recur, and consider how government should respond. The first dilemma is whether to ignore or rebut the theory; the second is whether to address the supply side of conspiracy theorizing by attempting to debias or disable its purveyors, to address the demand side by attempting to immunize third-party audiences from the theory's effects, or to do both (if resource constraints permit).

57. In both cases, the underlying structure of the problem is that conspiracy theorizing is a multi-party game. Government is faced with suppliers of conspiracy theories, and might aim at least in part to persuade, debias, or silence those suppliers. However, those two players are competing for the hearts and minds of third parties, especially the mass audience of the uncommitted.⁷⁸ Expanding the cast further, one may see the game as involving four players: government officials, conspiracy theorists, mass audiences, and independent experts—such as mainstream scientists—whom government attempts to enlist to give credibility to its rebuttal efforts.

i. Ignore or Rebut?

58. The first dilemma is that either ignoring or rebutting a conspiracy theory has distinctive risks and costs.⁷⁹ Ignoring the theory allows its proponents to draw ominous inferences from the government's silence. If the theory stands unrebutted, people may pay less attention to it, and even when they notice it, a natural inference from the government's silence is that the theory is too ludicrous to need rebuttal. But another possibility is that the government is silent because it cannot offer relevant evidence to the contrary. The suppliers of the conspiracy theories will propose the second inference. On this view, all misinformation (the initial conspiracy theory) should be met with counter misinformation.

59. On the other hand, to rebut the theory may be to legitimate it, moving the theory from the zone of claims too ludicrous to be discussed to the zone of claims that, whether or not true, are in some sense worth discussing. This legitimation effect can arise in one of two ways. First of all, third-party audiences may infer from the government's rebuttal efforts that the government itself estimates the conspiracy theory to be plausible and fears that the third parties will be persuaded. Here one risk is that the very act of rebuttal squarely focuses the audience on the conspiracy theory itself, in a way that may increase its salience and also its plausibility. Those who might reject the theory, or in any case not think about it, may take the rebuttal as a reason to give it serious consideration.

60. A second possible source of the legitimation effect is that some members of the audience may infer that many other members of the audience must believe the conspiracy theory, or government would not be taking the trouble to rebut it. Consider circumstances of "pluralistic ignorance," in which citizens are unsure what other citizens believe.⁸⁰ Citizens may take the fact of rebuttal itself as supplying information about the beliefs of other citizens, and may even use this information in forming their own beliefs. If the number who follow this cognitive strategy and thus adopt a belief in the conspiracy theory exceeds the number who are persuaded by the rebuttal, the perverse result of the rebuttal may then be to increase the number of believers.

61. In a typical pattern, government plays a wait-and-see strategy: ignore the conspiracy theory until it reaches some ill-defined threshold level of widespread popularity, and then rebut. There is a straightforward logic to this strategy. When the government ignores the theory, either the relevant audiences will draw an inference that the theory is silly, or else will infer that the government cannot effectively deny it. If the conspiracy theory does not spread despite the government's silence, the former inference is probably dominant, and response is unnecessary. There is also an option value to the strategy of ignoring the theory: a public rebuttal now is costly or impossible to undo, but maintaining silence now leaves government with the option to rebut later, if it chooses to do so. Finally and most generally, it seems silly and infeasible to chase after and rebut every conspiracy theory that comes to government's attention.

62. However, the concern that rebuttal will inadvertently legitimate a conspiracy theory overlooks an important synergistic gain: rebutting many conspiracy theories can reduce the legitimating effect of rebutting any one of them. When government rebuts a particular theory while ignoring most others, the legitimating effect arises at least in part because of a contrast between the foreground and the background: the inference is that government has picked the theory it is rebutting out of the larger set because this theory, unlike the others, is inherently plausible or is gaining traction among some sectors of the mass audience. The more theories government rebuts, the weaker is the implicit legitimating signal sent by the very fact of rebuttal.

63. It is impossible to say, in the abstract, how great this synergistic gain may be. However, the implication is that government should rebut more conspiracy theories than it would otherwise choose, if assessing the expected costs and benefits of rebuttal on a theory-by-theory basis. Because of synergy effects, government action considered over an array or range of cases may have different total costs and benefits than when those cases are considered one by one.

ii. Which Audience?

64. Should governmental responses be addressed to the suppliers, with a view to persuading or silencing them, or rather be addressed to the mass audience, with a view to inoculating them from pernicious theories? Of course these two strategies are not mutually exclusive; perhaps the best approach is to straddle the two audiences with a single response or simply to provide multiple responses. However, if there are resource constraints or intrinsic tradeoffs, such that the arguments that appeal to one audience tend to alienate the other, government will face a choice about what mix of responses is optimal

65. The basic problem with pitching governmental responses to the suppliers of conspiracy theories is that, as we have noted, those theories have a self-sealing quality. They are (1) resistant and in extreme cases invulnerable to contrary evidence,⁸¹ and (2) especially resistant to contrary evidence offered by the government, because the government rebuttal is folded into the conspiracy theory itself. If conspiracy theorists are responding to the informational signals given by those whom they trust, then the government's effort at rebuttal seems unlikely to be effective, and might serve to fortify rather than to undermine the original belief. After 9/11, one complex of conspiracy theories involved American Airlines Flight 77, which hijackers crashed into the Pentagon. Even those conspiracists who were persuaded that the Flight 77 conspiracy theories were wrong folded that view into a larger conspiracy theory. The problem with the theory that no plane hit the Pentagon, they said, is that the theory was too transparently false, disproved by multiple witnesses and much physical evidence. Thus the theory must have been a straw man initially planted by the government, in order to discredit other conspiracy theories and theorists by association.⁸²

66. Government can partially circumvent these problems if it enlists credible independent experts in the effort to rebut the theories. There is a tradeoff between credibility and control, however. The price of credibility is that government cannot be seen to control the independent experts. Although government can supply these independent experts with information and perhaps prod them into action from behind the scenes, too close a connection will prove self-defeating if it is exposed—as witnessed in the humiliating disclosures showing that apparently independent opinions on scientific and regulatory questions were in fact paid for by think-tanks with ties to the Bush administration.⁸³ Even apart from this tradeoff, conspiracy theorists may still fold independent third-party rebuttals into their theory by making conspiratorial claims of connection between the third party and the government. When the magazine *Popular Mechanics* offered a rebuttal of 9/11 conspiracy theories, conspiracists claimed that one of the magazine's reporters, Ben Chertoff, was the cousin of Homeland Security Secretary Michael Chertoff and was spreading disinformation at the latter's behest.⁸⁴

67. Because of these difficulties, many officials dismiss direct responses to the suppliers of conspiracy theorists as an exercise in futility. Those with strong commitments often engage in "biased assimilation" of evidence,⁸⁵ and conspiracy theorists are likely to be especially biased assimilators. Thus officials address their responses to the third-party mass audience, hoping to stem the spread of conspiracy theories by dampening the demand rather than by reducing the supply. When the National Institute of Standards and Technology issued a fact sheet to disprove the theory that the World Trade Center was brought down by a controlled demolition, the government spokesman stated that "[w]e realize this fact sheet won't convince those who hold to the alternative theories that our findings are sound. In fact, the fact sheet was never intended for them. It is for the masses who have seen or heard the alternative theory claims and want balance."⁸⁶

68. The problem with this line of argument, however, is that there are intrinsic costs to the strategy of giving up on the hard core of conspiracy theorists and directing government efforts solely towards inoculating the mass audience. For one thing, the hard core may itself provide the most serious threat. For another, a response geared to a mass audience (whether or not nominally pitched as a response to the conspiracy theorists) will lead some to embrace rather than reject the conspiracy theory the government is trying to rebut. This is the legitimization dilemma again: to begin a program of inoculation is to signal that the disease is already widespread and threatening. Under pluralistic ignorance, the perverse result may actually be to spread the conspiracy theory further.

iii. Cognitive Infiltration and Persuasion

69. Rather than taking the continued existence of the hard core as a constraint, and addressing itself solely to the third-party mass audience, government might undertake (legal) tactics for breaking up the tight cognitive clusters of extremist theories, arguments and rhetoric that are produced by the hard core and reinforce it in turn. One potentially promising tactic is cognitive infiltration of extremist groups. By this we do not mean 1960s-style infiltration with a view to surveillance and collecting information, possibly for use in future prosecutions. Rather, we mean that government efforts might succeed in weakening or even breaking up the epistemological complexes that constitute these networks and groups.

70. Recall that extremist networks and groups, including the groups that purvey conspiracy theories, typically suffer from a kind of crippled epistemology. We suggest a role for government efforts, and agents, in introducing cognitive diversity. Government agents (and their allies) might enter chat rooms, online social networks, or even real-space groups and attempt to undermine percolating conspiracy theories by raising doubts about their factual premises, causal logic, or implications for action, political or otherwise.

71. In one variant, government agents would openly proclaim, or at least make no effort to conceal, their institutional affiliations. A recent newspaper story recounts that Arabic-speaking Muslim officials from the State Department have participated in dialogues at radical Islamist chat rooms and websites in order to ventilate arguments not usually heard among the groups that cluster around those sites, with some success.⁸⁷ In another variant, government officials would participate anonymously or even with false identities. Each approach has distinct costs and benefits; the second risks perverse results but potentially brings higher returns.

72. In the former case, where government officials participate openly as such, hard-core members of the relevant networks, communities, and conspiracy-minded organizations may entirely discount what the officials say, right from the beginning. Because conspiracy theorists are likely to approach evidence and arguments in a biased way,⁸⁸ they are not likely to respond well, or even logically, to the claims of public officials. Of course government agents offer arguments and evidence against the conspiracy theory; perhaps their efforts are merely additional proof that the theory is correct. But the self-sealing quality of conspiracy theories (we should emphasize) is a matter of degree. Those who hold such theories may not be totally impervious to contrary evidence, even if it comes from those who are thought to have a stake in persuasion.

73. The risk with tactics of anonymous participation is that those tactics may be discovered or disclosed, with possibly perverse results. If the tactic becomes known, the conspiracy theory may become further entrenched, and any genuine member of the relevant groups who raises doubts may be suspected of government connections. And as we have emphasized throughout, in an open society it is difficult to conceal government conspiracies, even the sort of conspiratorial tactic we have suggested, whose aim is to undermine false and harmful conspiracy theorizing.

74. If disclosure of the tactic does occur, however, the perverse results are just a possible cost, whose risk and magnitude is unclear. Another possibility is that disclosure of the government's tactics will sow uncertainty and distrust within conspiratorial groups and among their members; new recruits will be suspect and participants in the group's virtual networks will doubt each other's bona fides. To the extent that these effects raise the costs of organization and communication for, and within, conspiratorial groups, the effects are desirable, not perverse. (And both sets of effects might occur simultaneously). So the two forms of cognitive infiltration offer evidently different risk-reward mixes. And despite the dangers, both are potentially useful instruments.

75. There is a similar tradeoff along another dimension: whether the infiltration should occur in the real world, through physical penetration of conspiracist groups by undercover agents, or instead should occur strictly in cyberspace. The latter is safer, but potentially less productive. The former will sometimes be indispensable, where the groups that purvey conspiracy theories (and perhaps themselves formulate conspiracies) formulate their views through real-space informational networks rather than virtual networks. Infiltration of any kind poses well-known risks: Perhaps agents will be asked to perform criminal acts to prove their bona fides, or (less plausibly) will themselves become persuaded by the conspiratorial views they are supposed to be undermining; perhaps agents will be unmasked and harmed by the infiltrated group. But the risks are generally greater for real-world infiltration, where the agent is exposed to more serious harms. Our main suggestion is just that, whatever the tactical details, there would seem to be ample reason for government efforts to introduce some cognitive diversity into the groups that generate conspiracy theories.⁸⁹

IV. Conclusion

Our goal here has been to understand the sources of false and harmful conspiracy theories and to examine potential government responses. Most people lack direct or personal information about the explanations for terrible events, and they are often tempted to attribute such events to some nefarious actor, in part because of their outrage. The temptation is least likely to be resisted if others are making the same attributions. Conspiracy cascades arise through the same processes that fuel many kinds of social errors. What makes such cascades most distinctive, and relevantly different from other cascades involving beliefs that are also both false and harmful, is their self-insulating quality. The very statements and facts that might dissolve conspiracy cascades can be taken as further evidence on their behalf. These points make it especially difficult for outsiders, including governments, to debunk them.

Some false conspiracy theories create serious risks. They do not merely undermine democratic debate; in extreme cases, they create or fuel violence. If government can dispel such theories, it should do so. One problem is that its efforts might be counterproductive, because efforts to rebut conspiracy theories also legitimate them. We have suggested, however, that government can minimize this effect

by rebutting more rather than fewer theories, by enlisting independent groups to supply rebuttals, and by cognitive infiltration designed to break up the crippled epistemology of conspiracy-minded groups and informationally isolated social networks.